

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

Search Results[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#) e-mail

Results for "(((cache<near/2>coherency)<and>(invalidat*<or>purg*<or>flush*<or>evict...")"

Your search matched 0 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance in Descending order**.» [Search Options](#)[View Session History](#)[Modify Search](#)[New Search](#)  Check to search only within this results setDisplay Format: Citation Citation & Abstract» [Key](#)

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STO IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistance.

[Help](#) [Contact Us](#) [Privacy & I](#)

© Copyright 2005 IEEE ...

Indexed by



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: [The ACM Digital Library](#) [The Guide](#)


THE ACM DIGITAL LIBRARY
 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used
[cache coherency](#) and [invalidat](#) or [purg](#) or [flush](#) or [evict](#) and [size](#) or [bucket](#) or [capacity](#) and [cache](#)

 Found
6,699 of
160,172

 Sort results by
 [Save results to a Binder](#)
[Try an Advanced Search](#)

 Display results
 [Search Tips](#)
[Try this search in The ACM Guide](#)
 [Open results in a new window](#)

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

 Relevance scale
1 Articles: Division of Labor in Embedded Systems


Ivan Goddard

 April 2003 **Queue**, Volume 1 Issue 2

 Full text available: [html\(37.05 KB\)](#) Additional Information: [full citation](#), [index terms](#)
2 Cache Memories


Alan Jay Smith

 September 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 3

 Full text available: [pdf\(4.61 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)
3 Separating data and control transfer in distributed operating systems


Chandramohan A. Thekkath, Henry M. Levy, Edward D. Lazowska

 November 1994 **Proceedings of the sixth international conference on Architectural support for programming languages and operating systems**, Volume 29 , 28 Issue 11 , 5

 Full text available: [pdf\(1.42 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Advances in processor architecture and technology have resulted in workstations in the 100+ MIPS range. As well, newer local-area networks such as ATM promise a ten- to hundred-fold increase in throughput, much reduced latency, greater scalability, and greatly increased reliability, when compared to current LANs such as Ethernet. We believe that these new network and processor technologies will permit tighter coupling of distributed systems at the hardware level, and that distribu ...

4 Caching considerations for generational garbage collection


Paul R. Wilson, Michael S. Lam, Thomas G. Moher

 January 1992 **ACM SIGPLAN Lisp Pointers , Proceedings of the 1992 ACM conference on LISP and functional programming**, Volume V Issue 1

 Full text available: [pdf\(1.09 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

5 Hardware-assisted replay of multiprocessor programs

David F. Bacon, Seth Copen Goldstein

December 1991 **ACM SIGPLAN Notices , Proceedings of the 1991 ACM/ONR workshop on Parallel and distributed debugging**, Volume 26 Issue 12

Full text available:  pdf(1.20 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



6 Architecture support for single address space operating systems

Eric J. Koldinger, Jeffrey S. Chase, Susan J. Eggers

September 1992 **ACM SIGPLAN Notices , Proceedings of the fifth international conference on Architectural support for programming languages and operating systems**, Volume 27 Issue 9

Full text available:  pdf(1.39 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



7 Increasing TLB reach using superpages backed by shadow memory.

Mark Swanson, Leigh Stoller, John Carter

April 1998 **ACM SIGARCH Computer Architecture News , Proceedings of the 25th annual international symposium on Computer architecture**, Volume 26 Issue 3

Full text available:  pdf(1.32 MB)  Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



The amount of memory that can be accessed without causing a TLB fault, the reach of a TLB, is failing to keep pace with the increasingly large working sets of applications. We propose to extend TLB reach via a novel Memory Controller TLB (MTLB) that lets us aggressively create superpages from non-contiguous, unaligned regions of physical memory. This flexibility increases the OS's ability to use superpages on arbitrary application data. The MTLB supports shadow pages, regions of physical address ...

8 Techniques for reducing overheads of shared-memory multiprocessing

Alain Kägi, Nagi Aboulenein, Douglas C. Burger, James R. Goodman

July 1995 **Proceedings of the 9th international conference on Supercomputing**

Full text available:  pdf(1.30 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



9 Analytical modelling of a hierarchical buffer for a data sharing environment

Asit Dan, Daniel M. Dias, Philip S. Yu

April 1991 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1991 ACM SIGMETRICS conference on Measurement and modeling of computer systems**, Volume 19 Issue 1

Full text available:  pdf(1.34 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



In a data sharing environment, where a number of loosely coupled computing nodes share a common storage subsystem, the effectiveness of a private buffer at each node is limited due to the multi-system invalidation effect, particularly under a non-uniform data access pattern. A global shared buffer can be introduced to alleviate this problem either as a disk cache or shared memory. In this paper we developed an approximate analytic model to evaluate different shared buffer management policies (SB ...)

10 Architectural primitives for a scalable shared memory multiprocessor

Joonwon Lee, Umakishore Ramachandran



June 1991 **Proceedings of the third annual ACM symposium on Parallel algorithms and architectures**

Full text available:  pdf(1.27 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

11 [WebExpress: a client/intercept based system for optimizing Web browsing in a wireless environment](#)



Barron C. Housel, George Samaras, David B. Lindquist

December 1998 **Mobile Networks and Applications**, Volume 3 Issue 4

Full text available:  pdf(338.35 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes an application model and software technology that makes it possible to run World Wide Web applications in wide area wireless networks. Web technology in conjunction with today's mobile devices (e.g., laptops, notebooks, personal digital assistants) and the emerging wireless technologies (e.g., digital cellular, packet radio, CDPD) offer the potential for unprecedented access to data and applications by mobile workers. Yet, the limited bandwidth, high latency, high cost, ...

12 [The evolution of Coda](#)



M. Satyanarayanan

May 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 2

Full text available:  pdf(441.35 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Failure-resilient, scalable, and secure read-write access to shared information by mobile and static users over wireless and wired networks is a fundamental computing challenge. In this article, we describe how the Coda file system has evolved to meet this challenge through the development of mechanisms for server replication, disconnected operation, adaptive use of weak connectivity, isolation-only transactions, translucent caching, and opportunistic exploitation of hardware surrogates. For eac ...

Keywords: Adaptation, Linux, UNIX, Windows, caching, conflict resolution, continuous data access, data staging, disaster recovery, disconnected operation, failure, high availability, hoarding, intermittent networks, isolation-only transactions, low-bandwidth networks, mobile computing, optimistic replica control, server replication, translucent cache management, weakly connected operation

13 [Coherency for multiprocessor virtual address caches](#)



James R. Goodman

October 1987 **Proceedings of the second international conference on Architectural support for programming languages and operating systems**, Volume 15 , 22 , 21 Issue 5 , 10 , 4

Full text available:  pdf(962.36 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A multiprocessor cache memory system is described that supplies data to the processor based on virtual addresses, but maintains consistency in the main memory, both across caches and across virtual address spaces. Pages in the same or different address spaces may be mapped to share a single physical page. The same hardware is used for maintaining consistency both among caches and among virtual addresses. Three different notions of a cache "block" are defined: (1) the unit for transferring data t ...

14 1 - Regular Articles: Cache-conscious sorting of large sets of strings with dynamic tries



Ranjan Sinha, Justin Zobel

December 2004 **Journal of Experimental Algorithms (JEA)**, Volume 9

Full text available:  pdf(848.84 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Ongoing changes in computer architecture are affecting the efficiency of string-sorting algorithms. The size of main memory in typical computers continues to grow but memory accesses require increasing numbers of instruction cycles, which is a problem for the most efficient of the existing string-sorting algorithms as they do not utilize cache well for large data sets. We propose a new sorting algorithm for strings, burstsort, based on dynamic construction of a compact trie in which strings are ...

15 Cache memory performance in a unix environment



Cedell Alexander, William Keshlare, Furrokh Cooper, Faye Briggs

June 1986 **ACM SIGARCH Computer Architecture News**, Volume 14 Issue 3

Full text available:  pdf(2.10 MB) Additional Information: [full citation](#), [citations](#), [index terms](#)

16 Coarse-grained parallelism for hierarchical radiosity using group iterative methods



Thomas A. Funkhouser

August 1996 **Proceedings of the 23rd annual conference on Computer graphics and interactive techniques**

Full text available:  pdf(358.69 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

17 Cache performance for multimedia applications



Nathan T. Slingerland, Alan Jay Smith

June 2001 **Proceedings of the 15th international conference on Supercomputing**

Full text available:  pdf(642.63 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The caching behavior of multimedia applications has been described as having high instruction reference locality within small loops, very large working sets, and poor data cache performance due to non-locality of data references. Despite this, there is no published research deriving or measuring these qualities. Utilizing the previously developed Berkeley Multimedia Workload, we present the results of execution driven cache simulations with the goal of aiding future media processing architect ...

Keywords: CPU caches, cache, mulitmedia, simulation, trace driven simulation

18 An analytical cache model



A. Agarwal, J. Hennessy, M. Horowitz

May 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 2

Full text available:  pdf(2.51 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Trace-driven simulation and hardware measurement are the techniques most often used to obtain accurate performance figures for caches. The former requires a large amount of simulation time to evaluate each cache configuration while the latter is restricted to measurements of existing caches. An analytical cache model that uses parameters extracted from address traces of programs can efficiently provide estimates of cache performance and show the effects of varying cache parameters. By repre ...

19 Empirical performance evaluation of concurrency and coherency control protocols for database sharing systems



Erhard Rahm

June 1993 **ACM Transactions on Database Systems (TODS)**, Volume 18 Issue 2

Full text available:  pdf(3.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database Sharing (DB-sharing) refers to a general approach for building a distributed high performance transaction system. The nodes of a DB-sharing system are locally coupled via a high-speed interconnect and share a common database at the disk level. This is also known as a "shared disk" approach. We compare database sharing with the database partitioning (shared nothing) approach and discuss the functional DBMS components that require new and coordinated solutions for DB-shar ...

Keywords: coherency control, concurrency control, database partitioning, database sharing, performance analysis, shared disk, shared nothing, trace-driven simulation

20 Hierarchical cache/bus architecture for shared memory multiprocessors 

A. W. Wilson

June 1987 **Proceedings of the 14th annual international symposium on Computer architecture**

Full text available:  pdf(978.39 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A new, large scale multiprocessor architecture is presented in this paper. The architecture consists of hierarchies of shared buses and caches. Extended versions of shared bus multicache coherency protocols are used to maintain coherency among all caches in the system. After explaining the basic operation of the strict hierarchical approach, a clustered system is introduced which distributes the memory among groups of processors. Results of simulations are presented which demonstrate that t ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  Adobe Acrobat  QuickTime  Windows Media Player  Real Player

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	3901	cache adj coherency	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:44
L2	2315	(invalidat\$3 or purg\$ or flush\$ or evict\$4) and 1	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:54
L3	14221	cache with (size or bucket or capacity)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:46
L4	1214	2 and 3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:46
L5	4560	delay with (initializ\$ or stat\$1up)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:47
L6	1	4 and 5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:47
L7	9	vmscluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:48
L8	8	7 and 4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:48
L9	52	vaxcluster	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:48

L10	8	4 and 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:49
L11	174	cache adj software	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:49
L12	30	4 and 11	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:52
L13	28	12 and ((i/o adj device) or storage or disk)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:53
L14	2018	((invalidat\$3 or purg\$ or flush\$ or evict\$4) with cach\$4) and 1	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:54
L15	1103	14 and 3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:55
L16	1	5 and 15	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/08/18 21:55